

Improving Scalability of Reinforcement Learning by Separation of Concerns

Harm van Seijen¹

harm.vanseijen@maluuba.com

Mehdi Fatemi¹

mehdi.fatemi@maluuba.com

Joshua Romoff^{1,2}

joshua.romoff@mail.mcgill.ca

¹Maluuba Research, Montreal, Canada

²McGill University, Montreal, Canada

Abstract

In this paper, we propose a framework for solving a single-agent task by using multiple agents, each focusing on different aspects of the task. This approach has two main advantages: 1) it allows for specialized agents for different parts of the task, and 2) it provides a new way to transfer knowledge, by transferring trained agents. Our framework generalizes the traditional hierarchical decomposition, in which, at any moment in time, a single agent has control until it has solved its particular subtask. We illustrate our framework using a number of examples.

1 Introduction

Marvin Minsky's "Society of Mind" theory postulates that our behaviour is not the result of a single cognitive agent, but rather the result of a society of individually simple, interacting processes called agents (Minsky, 1988). The power of this approach lies in specialization: different agents can have different representations, different learning processes, and so on. On a larger scale, our society as a whole validates this approach: our technological achievements are the result of many cooperating specialized agents.

In reinforcement learning (RL), where the goal is to learn a policy for an agent interacting with an initially unknown environment, the importance of breaking large tasks into smaller pieces has long been recognized. Specifically, there has been a lot of work on hierarchical RL methods, which decompose a task in a hierarchical way into subtasks (Dietterich, 2000; Barto & Mahadevan, 2003; Kulkarni et al., 2016). Hierarchical learning can help accelerate learning on individual tasks by mitigating the exploration challenge of sparse-reward problems. One of the most popular frameworks for this is the options framework (Sutton et al., 1999), which extends the standard RL framework based on Markov Decision processes (MDP) with temporally-extended actions.

In this paper, we propose a framework of communicating agents that aims to generalize the traditional hierarchical decomposition to allow for more flexible task decompositions. For example, decompositions where multiple subtasks have to be solved in parallel, or in cases where a subtask does not have a well-defined end, but rather is a continuing process that needs constant adjustment, such as walking through a crowded street. We refer to our framework as separation-of-concerns.

What is unique about our proposed framework is the way agents cooperate. To enable cooperation, the reward function of a specific agent not only has a component that depends on the environment state, but also a component that depends on the communication actions of the other agents. Depending on the specific mixture of these components, agents have different degrees of independence. In

addition, because the reward in general is state-specific, an agent can show different levels of dependence in different parts of the state-space. Typically, in areas with high environment-reward, an agent will act independent of the communication actions of other agents, while in areas with low environment-reward, an agent’s policy will depend strongly on the communication actions.

In this initial work, we focus primarily on the theoretical aspects of our framework. Formally, our framework can be seen as a sequential multi-agent decision making system with non-cooperative agents. This is a challenging setting, because from the perspective of one agent, the environment is non-stationary due to the learning of the other agents. We address this by defining trainer agents with a fixed policy. Learning with these trainer agents can occur in two ways: 1) by pre-training agents and then freezing their policy, or 2) by learning in parallel using off-policy learning.

The rest of this document is organized as follows. In Section 2, we discuss related work. In Section 3, we introduce our model. Section 4 contains a number of experiments. Section 5 discusses the results and future work. Finally, section 6 concludes.

2 Related Work

Hierarchical Learning / Options Our work builds upon the long line of work on hierarchical learning. Dietterich (2000) introduced the MAXQ framework, which decomposes the value function in a hierarchical way. Sutton et al. (1999) introduced the options framework. Options are temporally extended actions consisting of an initialization set, an option policy and a termination condition. Effectively, applying options to an MDP, changes it into a semi-MDP. Options have also been popular as a mechanism for skill discovery.

There has been significant work on option discovery. In the tabular setting, useful subgoal states can be identified, for example, by using heuristics based on the visitation frequency (McGovern & Barto, 2001), by using graph partitioning techniques (Menache et al., 2002; Şimşek et al., 2005), or by using the frequency with which state variables change (Hengst, 2002). With function approximation, finding good subgoals becomes significantly more challenging. Konidaris & Barto (2009) assumed that subgoal states were given (hence, only the option policy needs to be learned). Machado & Bowling (2016) perform option discovery by identifying ‘purposes’ at the edge of a random agent’s visitation area. Learning options towards those edge-purposes brings the agent quickly to a new region where it can continue exploration. Bacon et al. (2016) propose a new architecture to learn the policy over options, the options themselves, as well as their respective termination conditions. This is accomplished without defining any particular subgoal. Only the number of options is known beforehand.

Kulkarni et al. (2016) study hierarchical RL in the context of deep reinforcement learning. In their setting, a high-level controller specifies a goal for the low-level controller. Once the goal is accomplished, the top-level controller selects a new goal for the low-level controller. The system was trained in two phases: in the first phase the low-level controller is trained on a set of different goals; in the second phase the high-level and low-level controllers are trained in parallel. Heess et al. (2016) also use a system with a high-level and a low-level controller, but the high-level controller continuously sends a modulation signal to the low-level controller, affecting its policy. This setup can be seen as a special case of the framework that we introduce.

Multi-agent systems There is also a large amount of work on multi-agent RL (e.g., see Buşoniu et al. (2010) for an overview). The standard multi-agent configuration includes multiple agents which are acting on the environment simultaneously and which receive rewards individually based on the joint actions. This can naturally be modelled as a stochastic game.

Multi-agent systems can be divided into fully cooperative, fully competitive or mixed tasks (neither cooperative nor competitive). For a fully cooperative task, all agents share the same reward function. Early works includes the Integrated Learning System (ILS) by Silver et al. (1990), which integrates heterogeneous learning agents (such as search-based and knowledge-based) under a central controller through which the agents critique each other’s proposals. Alternatively, Whitehead (1991) proposed Learning with an External Critic (LEC) and Learning by Watching (LBW), which advocate learning from other agents in a social settings. It has been shown that a society of n Q-learning agents, which are watching each other can learn $O(\frac{1}{n})$ faster than a single Q-learning agent. More recently, Foerster et al. (2016) used a framework of communicating agents based on deep neural networks to solve

various complex tasks. They evaluated two learning approaches. For the first approach, each agent learns its own network parameters, while treating the other agents as part of the environment. The second approach uses centralized learning and passes gradients between agents.

For fully competitive tasks, which typically deal with the two-agent case only, the agents have opposing goals (the reward function of one agent is the negative of the reward function of the other). Our work falls into the mixed category, which does not pose any restrictions on the reward functions. A lot of work in the mixed setting focuses on static (stateless) tasks.

Separation of Concerns (SoC) can benefit from ideas and theoretical advances in multi-agent frameworks. Nevertheless, it should be noted that there are key differences between SoC and multi-agent methodologies. Intrinsically, SoC splits a single-agent problem into multiple parallel, communicating agents with simpler and more focused, but different objectives (skills).

3 Separation of Concerns Model

In this section, we introduce our model for dividing a single-agent task into a multi-agent task.

3.1 The Model

Our single-agent task is defined by a Markov decision process (MDP), consisting of the tuple $\langle \mathcal{X}, \mathcal{A}, p, r, \gamma \rangle$, where \mathcal{X} is the set of states; \mathcal{A} the set of actions; $p(s'|s, a)$ indicates the probability of a transition to state $s' \in \mathcal{S}$, when action $a \in \mathcal{A}$ is taken in state $s \in \mathcal{S}$; $r(s, a, s')$ indicates the reward for a transition from state s to state s' under action a ; finally, the discount factor γ specifies how future rewards are weighted with respect to the immediate reward. Actions are taken at discrete time steps according to policy π , which maps states to actions. The goal is to maximize the discounted sum of rewards, also referred to as the *return*: $G_t := \sum_{k=1}^{\infty} \gamma^{k-1} R_{t+k}$. We call the agent of the single-agent task the *flat agent*.

We expand this task into a system of n communicating agents as follows. For each agent i , we define an environment action-set \mathcal{B}^i , a communication action-set \mathcal{C}^i , and a *learning objective*, defined by a reward function r^i plus a discount factor γ^i . Furthermore, we define an action-mapping function $f : \mathcal{B}^1 \times \dots \mathcal{B}^n \rightarrow \mathcal{A}$ that maps the joint environment-action space to an action of the flat agent. The agents share a common state-space \mathcal{Y} consisting of the state-space of the flat agent plus the joint communication actions: $\mathcal{Y} := \mathcal{X} \times \mathcal{C}^1 \times \dots \times \mathcal{C}^n$. See Figure 1 for an illustration for $n = 2$.

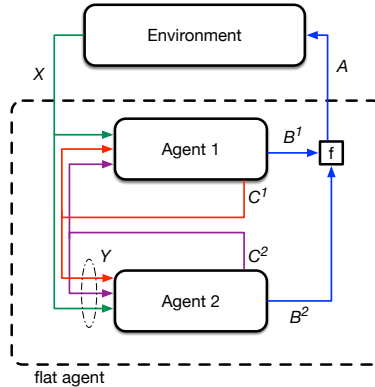


Figure 1: SoC model for 2 agents

At time t , each agent i observes state $Y_t := (X_t, C_{t-1}^1, \dots, C_{t-1}^n) \in \mathcal{Y}$ and selects environment action B_t^i and communication action $C_t^i \in \mathcal{C}^i$ according to policy $\pi^i : \mathcal{Y} \rightarrow \mathcal{B}^i \times \mathcal{C}^i$. Action $A_t = f(B_t^1, \dots, B_t^n)$ is fed to the environment, which responds with an updated state X_{t+1} . The environment also produces a reward R_{t+1} , but this reward is only used to measure the overall performance of the SoC model. For learning, each agent i uses its own reward function $r^i : \mathcal{Y} \times \mathcal{B}^i \times \mathcal{C}^i \times \mathcal{Y} \rightarrow \mathbb{R}$ to compute reward $R_{t+1}^i = r^i(Y_t, B_t^i, C_t^i, Y_{t+1})$.

An important property of the SoC model is that the reward function of a particular agent depends on the communication actions of the other agents. This provides a clear incentive for an agent to react in response to communication, even in the case of full observability. For example, agent A can ‘ask’ agent B to behave in a certain way via a communication action that rewards agent B for this behaviour.

Full observability is not an explicit requirement of our framework. The general model described above can be extended in different ways. In particular, extensions that allow for further specialization of agents will increase the benefit of the SoC model as a whole. Some examples are:

- state abstractions/information hiding: because the agents have different objectives, they can use different state-abstractions.
- action-selection at different frequencies.
- state-dependent gamma (such that terminal states can be defined)
- state-dependent action sets.

Example: Hierarchical Configuration

One important example is the case where the agents are organized in a way that decomposes a task hierarchically. As an example, consider a configuration consisting of three agents: agent 0 is the top-level agent, and agent 1 and agent 2 are two bottom-level agents. The top-level agent only has communication actions, specifying which of the bottom level agents is in control. That is, $\mathcal{B}^0 = \emptyset$ and $\mathcal{C}^0 = \{\text{‘agent 1 in control’}, \text{‘agent 2 in control’}\}$. Agent 1 and agent 2 both have a state-dependent action-set that gives access to the environment actions \mathcal{A} if they have been given control by agent 0. That is, for agent 1:

$$\mathcal{B}^1 = \begin{cases} \mathcal{A} & \text{if } c^0 = \text{‘agent 1 in control’} \\ \text{‘none’} & \text{if } c^0 = \text{‘agent 2 in control’} \end{cases},$$

and vice-versa for agent 2. By allowing agent 0 to only switch its action once the agent currently in control has reached a terminal state (by either storing a set of terminal state conditions itself, or by being informed via a communication action), a typical hierarchical task decomposition is achieved.

This example illustrates that our SoC model is a generalization of a hierarchical model. Like any multi-agent system, obtaining stable performance for some SoC configurations can be challenging. In the next section, we discuss some learning strategies for obtaining stability.

3.2 Learning Strategies

A common learning approach for mixed-strategy multi-agent systems like ours is to use a single-agent algorithm for each of the agents. While there is no convergence guarantee in this case, it often works well in practice (Shoham & Leyton-Brown, 2009). In this section, we examine under which conditions exact convergence can be guaranteed.

Proposition 1 *By assigning a stationary policy to all agents, except agent i , an implicit MDP is defined for agent i with state space \mathcal{Y} , reward function r^i and (joint) action space $\mathcal{B}^i \times \mathcal{C}^i$.*

Proof The proposition holds if the next state $Y_{t+1} := (X_{t+1}, C_t^1, \dots, C_t^n)$ only depends on the current state Y_t and joint action (B_t^i, C_t^i) . Because the policies of all agents other than agent i are fixed, knowing Y_t fixes a distribution over the environment and communication actions for each of the other agents. The distribution over these environment actions, together with the environment action of agent i determines a distribution for the random variable X_{t+1} . Together with the (distribution over) communication actions this fixes a distribution for Y_{t+1} .

It follows from Proposition 1 that if we also define a policy for agent i , then we get a well-defined value-function. Let η be a tuple of policies, assigning a policy to each agent: $\eta := \langle \pi^1, \dots, \pi^n \rangle$. We can now define a value-function $v^{i,\eta}$ with respect to reward function and discount factor of agent i as follows:

$$v^{i,\eta}(y) := \mathbb{E}\{R_{t+1}^i + \gamma R_{t+2}^i + \dots | Y_t = y, \eta\}$$

Using this, we define the following independence relation between agents.

Definition Agent i is independent of agent j if the value $v^{i,\eta}(y)$ does not depend on the policy of agent j . Agent i is dependent of agent j if it is not independent of agent j .

Note that the agent-state y includes the communication action of agent j . Hence, this independence definition still allows for a dependence on the most recent communication action of agent j ; the independence is only with respect to the future actions of agent j , that is, its policy. A simple example of a case where this independence relation holds is the hierarchical case, where the actions of the top agent remain fixed until the bottom agent reaches a terminal state.




Figure 2: Dependency graphs. An arrow pointing from agent j to agent i means that agent i depends on agent j . Circles represent regular agents; diamonds represent trainer agents. The cyclic graph shown in (b) can be transformed into an acyclic graph by adding trainer agents. A trainer agent for agent i defines fixed behaviour for the agents that agent i depends on to ensure stable learning.

Using the independence definition above, we can draw a *dependency graph* of an SoC system, which is a directed graph with arrows showing the dependence relations: an arrow pointing from agent j to agent i means that agent i depends on agent j . In general, a dependency graph can be acyclic (containing no directed cycles, see Figure 2a) or cyclic (containing directed cycles, see Figure 2b).

If the dependency graph is an acyclic graph, using single-agent Q-learning to train the different agents is especially straightforward, as shown by the following theorem.

Theorem 1 *When the dependency graph is acyclic, using Q-learning for each agent in parallel yields convergence under the standard (single-agent) conditions.*

Proof Because the graph is acyclic, some agents are independent of any of the other agents and hence will converge over time. Once the Q-values of these independent agents are sufficiently close to their optimal values (such that the distance to the optimal value is smaller than the difference between two actions), their policies will not change anymore. Once this occurs, the agents that only depend on these independent agents will converge. After this, the agents one step downstream in the graph will converge, and so on, until all agents have converged. ■

In many cases, the dependency graph will be cyclic. In this case, convergence of Q-learning is not guaranteed. However, we can transform a cyclic graph into an acyclic one by using *trainer agents*.

Definition A *trainer agent*, assigned to a particular agent i , is a fixed-policy agent that generates behaviour for the agents that agent i depends on.

Assigning a trainer agent to agent i implicitly defines a stationary MDP for agent i with a corresponding optimal policy that can be learned. Hence, agent i only depends on the trainer agent. The trainer agent itself is an independent agent. Hence, trainer agents can be used to break cycles in dependency graphs. Note that a cyclic graph can be transformed into an acyclic one in different ways. In practice, which agents are assigned trainer agents is a design choice that depends on how easy it is to define effective trainer behaviour. In the simplest case, a trainer agent can just be a random or semi-random policy. As an example, the graph shown in Figure 2b can be transformed into the graph shown in Figure 2c by defining a trainer agent $1'$ for agent 1 that generates behaviour for agent 2 and 3, and a trainer agent $2'$ for agent 2 that generates behaviour for agent 1 and 3.

Learning with trainer agents can occur in two ways. The easiest way is to pre-train agents with their respective trainer agents, then freeze their weights and train the rest of the agents. Alternatively, all agents could be learned in parallel, but the agents that are connected to a trainer agent use off-policy learning to learn values that correspond to the policy of the trainer agent, while the behaviour policy is generated by the regular agents. Off-policy learning can be achieved by importance sampling,

which corrects for the frequency at which a particular sample is observed under the behaviour policy versus the frequency at which it is observed under the target policy. This can be achieved as follows. Consider agent i with actions $\{a_1^i, a_2^i\}$ that depends on agent j with actions $\{a_1^j, a_2^j\}$. Furthermore, consider that agent i has a trainer agent i' attached to it mimicking behaviour for agent j . In other words, agent i' also has actions $\{a_1^j, a_2^j\}$. At any moment in time, the actual behaviour is generated by agents i and j . If at time t , agent j selects action a_1^j , while the selection probability for that action is w_t^j , and the selection probability for that same action is $w_t^{i'}$ for trainer agent i' , then the off-policy update for agent i is:

$$Q(Y_t, A_t^i) \leftarrow Q(Y_t, A_t^i) + \frac{w_t^{i'}}{w_t^j} \alpha (R_{t+1}^i + \max_a Q(Y_{t+1}, a) - Q(Y_t, A_t^i)).$$

Obtaining convergence does not necessarily mean that the policy that is obtained is a good policy. In the next section, we address the optimality of the policy.

3.3 Optimality

In the context of hierarchical learning, Dietterich (2000) defines *recursive optimality* as a type of local optimality, in which the policy for each subtask is optimal given the policies of its children-subtasks. The recursive optimal policy is the overall policy that consists of the combination of all the locally optimal policies. The recursive optimal policy is in general worse than the optimal policy for a flat agent, but it can be much easier to find.

We can define a similar form of optimality for an SoC model:

Definition If the dependency graph of an SoC model is acyclic (with or without adding trainer agents), then we define a recursive optimal SoC policy $\eta := \langle \pi^1, \dots, \pi^n \rangle$ as the policy consisting of all locally optimal policies, that is, policy π^i is optimal for agent i , given the policies of the agents that it depends on.

The learning strategies discussed in the previous section will converge to the recursive optimal policy. How close this policy is to the optimal policy depends on the specific decomposition, that is, the communication actions, the agent-reward functions and, potentially, the employed trainer agents.

4 Experiments

In this section, we illustrate the application of the SoC model using two examples. First, we demonstrate the method on a tabular domain. Then, we consider a pixel-based game and combine SoC with deep reinforcement learning.

4.1 Navigation Domain

The aim of the first experiment is to show the scalability of the SoC model. For this, we use a navigation domain, shown in Figure 3. The action set consists of a move forward action, a turn clockwise and a turn counterclockwise action. Furthermore, we add a varying number of extra ‘no-op’ actions (actions without effect) to control the complexity of the domain. The reward is -5 when the agent bumps into a wall and -1 for all other transitions.

We compare a flat agent with an SoC configuration consisting of a high and low-level agent. The high-level agent communicates a compass direction to the low-level agent, $\mathcal{C}^h = \{\text{north, south, east, west}\}$, and has no environmental actions ($\mathcal{B}^h = \emptyset$). The low-level agent in turn has access to all environmental actions $\mathcal{B}^l = \mathcal{A} = \{\text{move, clockwise turn, counter-clockwise turn, no-op1, no-op2, ...}\}$ and no communication actions ($\mathcal{C}^l = \emptyset$). The reward function of the high-level agent is such that it gets -1 on each transition. The reward function of the low-level agent is such that it gets -5 for hitting the wall and +1 if it makes a move in the direction requested by the high-level agent. All agents are trained with Q-learning and use ϵ -greedy exploration with a fixed ϵ of 0.01 and a step-size of 0.1.

The left of Figure 4 shows the learning behaviour for tasks with different levels of complexity. Specifically, we compare tasks with 5, 10 and 20 no-op actions. While the complexity has only a




Figure 3: Navigation task: the agent (triangular yellow object) has to navigate from the start square (in upper-left corner) to the goal square (in lower-right corner) using the three actions: move one block forward, turn 90° clockwise and turn 90° counter-clockwise. Furthermore, additional ‘no-op’ actions are added to increase the complexity of the learning task.




Figure 4: *[left]* Comparison of the SoC model and a flat agent in the navigation domain for three different task complexities. *[right]* Scalability of SoC learning: the average return over the first 4000 episodes on the navigation task as function of the number of added no-op actions.

small affect on the performance of the SoC method, it affects the flat agent considerably. This is further illustrated by the right graph of Figure 4, which shows the average return over 4000 episodes for a different number of no-op actions.

Overall, these results clearly illustrate the ability of the SoC model to improve the scalability. Note that implementing this with a hierarchical approach would require the high-level agent to know the available compass directions in each grid-cell to avoid giving the low-level agent a goal that it cannot fulfill (e.g., it cannot move ‘north’ while it is in the top-left corner). By contrast, the high-level agent of the SoC system does not require this information.

4.2 Catch

In our second example, we compare a flat agent with the SoC model on the game Catch. Catch is a simple pixel-based game introduced by Mnih et al. (2014). The game consists of a 24 by 24 screen of binary pixels in which the goal is to catch a ball that is dropped at a random location at the top of the screen with a paddle that moves along the bottom of the screen. In our case, both the ball and the paddle consist of just a single pixel. The available actions are `left`, `right` and `stay`. The agent receives +1 reward for catching the ball, -1 if the ball falls off the screen and 0 otherwise.

Similar to the navigation domain, our SoC model consists of a high-level and a low-level agent. The high-level agent has no direct access to the environment actions, but it communicates a desired action to the low-level agent: $\mathcal{C}^h = \{\text{left}, \text{right}, \text{stay}\}$. The low-level agent has direct access to the




Figure 5: Illustration of the network used for the flat agent and the high-level agent [a] versus the network used for the low-level agent [b]. Because the low-level agent uses a 10 by 10 bounding box, it does not require a full convolution network.

environment actions and no communication actions: $\mathcal{B}^l = \mathcal{A}$ and $\mathcal{C}^l = \emptyset$. Furthermore, the high-level agent has a discount factor of 0.99 and access to the full screen, whereas the low-level agent has a discount factor of 0.65 and uses a bounding box of 10 by 10 pixels around the paddle.¹ The low-level agent only observes the ball when it is inside the bounding box. The high-level agent receives a reward of +1 if the ball is caught and -1 otherwise; the low-level agent receives the same reward plus a small positive reward for taking the action suggested by the high-level agent. The high-level agent takes actions every 2 time steps, whereas the low-level agent takes actions every time step.

Both the flat agent and the high-level and low-level agents are trained using DQN (Mnih et al., 2015). The flat agent uses a convolutional neural network defined as follows. The 24x24 binary image is passed through two convolutional layers, followed by two dense layers. Both convolutional layers have 32 filters of size (5,5) and a stride of (2,2). The first dense layer has 128 units, followed by the output layer with 3 units. The network uses the same activations and initializations as in (Mnih et al., 2015).

The high-level agent in the SoC system uses an identical architecture to that of the flat agent. However, due to the reduced state size for the low-level agent, it only requires a small dense network. The network flattens the 10x10 input and passes it through two dense layers with 128 units each. The output is then concatenated with a 1-hot vector representing the communication action of the high-level agent. The merged output is then passed through one final dense layer with 3 units (see Figure 5).

The left graph of Figure 6 shows the results of the comparison. The SoC model learns significantly faster than the flat agent. To show the importance of the co-operation between the low-level and the high-level agent, we performed an additional experiment where we varied the additional reward the low-level agent gets for taking the action suggested by the high-level agent. The results are shown in the right graph of Figure 6. If the additional reward is 0, the low-level agent has no incentive to listen to the high-level agent and will act fully independent; alternatively, if the additional reward is very high, it will always follow the suggestion of the high-level agent. Because both agents are limited (the high-level agent has a low action-selection frequency, while the low-level agent has a limited view), both these situations are undesirable. The ideal low-level agent is one that acts neither fully independent nor fully dependent.

5 Discussion and Future Work

The experiments from the previous section showed the validity of the separation of concerns principle: separating a task into multiple related sub-tasks can result in considerable speed-ups. In the presented experiments, the decomposition was made a priori. In future work, we would like to focus on (partially) learning the decomposition. In this case, we do not necessarily expect an advantage on single tasks, due to the cost of learning the decomposition. In the transfer learning setting, however,

¹It works also without a bounding box, but we found that performance was better with a bounding box.




Figure 6: [left] Learning speed comparison on Catch (one epoch corresponds with 40 episodes). [right] Effect of the communication reward on the final performance of the SoC system.

where a high initial cost for learning a representation can be offset by many future applications of that representation, it could prove to be useful.

The SoC configuration we used in both our examples consisted of a high-level agent that only communicates and a low-level agent that only performs environment actions. Another direction for future work is to explore alternative configurations and use more than two agents.

The reward function in reinforcement learning often plays a double role: it acts as both the performance objective, specifying what type of behaviour is desired, as well as the learning objective, that is, the feedback signal that modifies the agent’s behaviour. That these two roles do not always combine well into a single function becomes clear from domains with sparse rewards, where learning can be prohibitively slow. Work on intrinsic motivation aims to mitigate this issue by providing an additional reward function to steer learning. In a sense, the SoC model takes this a step further: the performance objective, consisting of the reward function of the environment, is fully separated from the learning objective of the agents, consisting of the agent’s reward function.

This clear separation between performance objective and learning objective further separates the SoC model from options. Options, once learned, aggregate the rewards obtained from the environment. Hence, the top-level agent of a hierarchical system based on options learns a value function based on the environment reward. We argue for a clearer separation.

6 Conclusions

We presented initial work on a framework for solving single-agent tasks using multiple agents. In our framework, different agents are concerned with different parts of the task. Our framework can be viewed as a generalization of the traditional hierarchical decomposition. We identified conditions under which convergence of Q-learning occurs (to a recursive optimal policy) and showed experiments to validate the approach.

References

- Bacon, P., Harb, J., and Precup, D. The option-critic architecture, 2016. arXiv:1609.05140 [cs.AI].
- Barto, A. G. and Mahadevan, S. Recent advances in hierarchical reinforcement learning. *Discrete Event Dynamic Systems*, 13(4):341–379, 2003.
- Buşoniu, L., Babuška, R., and de Schutter, B. Multi-agent reinforcement learning: An overview. In *Innovations in Multi-Agent Systems and Applications-1*, pp. 183–221. Springer Berlin Heidelberg, 2010.
- Şimşek, Ö., Wolfe, A. P., and Barto, A. G. Identifying useful subgoals in reinforcement learning by local graph partitioning. In *In Proceedings of the International Conference on Machine Learning (ICML)*, 2005.

- Dietterich, T. G. Hierarchical reinforcement learning with the maxq value function decomposition. *Journal of Artificial Intelligence Research*, 13:227–303, 2000.
- Foerster, J. N., Assael, Y. M., de Freitas, N., and Whiteson, S. Learning to communicate with deep multi-agent reinforcement learning. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2016.
- Heess, N., Wayne, G., Tassa, Y., Lillicrap, T., Riedmiller, M., and Silver, D. Learning and transfer of modulated locomotor controllers, 2016. arXiv:1610.05182 [cs.RO].
- Hengst, B. Discovering hierarchy in reinforcement learning with HEXQ. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2002.
- Konidaris, G. and Barto, A. G. Skill discovery in continuous reinforcement learning domains using skill chaining. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2009.
- Kulkarni, T. D., Narasimhan, K. R., Saeedi, A., and Tenenbaum, J. B. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation, 2016. arXiv:1604.06057 [cs.LG].
- Machado, M. C. and Bowling, M. Learning purposeful behaviour in the absence of rewards, 2016. arXiv:1605.07700 [cs.LG].
- McGovern, A. and Barto, A. G. Automatic discovery of subgoals in reinforcement learning using diverse density. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2001.
- Menache, I., Mannor, S., and Shimkin, N. Q-cut - dynamic discovery of sub-goals in reinforcement learning. In *Proceedings of the European Conference on Machine Learning (ECML)*, 2002.
- Minsky, M. *Society of mind*. Simon and Schuster, 1988.
- Mnih, V., Heess, N., Graves, A., and Kavukcuoglu, K. Recurrent models of visual attention. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2014.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., Kumaran, H., King D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- Shoham, Y. and Leyton-Brown, K. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.
- Silver, B., Frawley, W., Iba, G., Vittal, J., and Bradford, K. ILS: A framework for multi-paradigmatic learning. In *Proceedings of the Seventh International Conference on Machine Learning (ICML)*, pp. 348–356, 1990.
- Sutton, R.S., Precup, D., and Singh, S.P. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2):181–211, 1999.
- Whitehead, S. D. A complexity analysis of cooperative mechanisms in reinforcement learning. In *Proceedings of AAAI-91*, pp. 607–613, 1991.